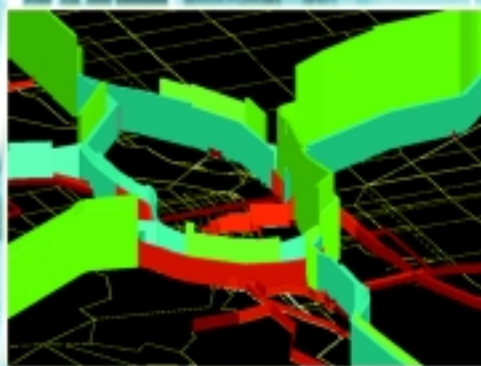
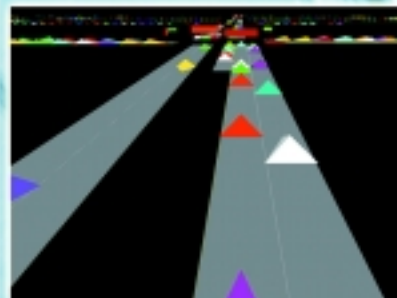


BITS

computing & communications news



Los Alamos
NATIONAL LABORATORY

June/July 2000

TRANSIMS—TRansportation ANalysis and SIMulation System

Revolutionizing Traffic Forecasting and Planning



A Laboratory team has created TRANSIMS, an extensive software system that simulates human mobility on an urban regional scale. The team was awarded Distinguished Copyright Award at the Laboratory's Patent and Licensing Awards Ceremony on March 1, 2000. TRANSIMS team members present at the Awards Ceremony are shown from left to right: Doug Roberts (TSA-2), Mike Williams (TSA-4), Madhav Marathe (TSA-2), Phil Romero (CIC-12), Deborah Kubicek (TSA-5),

Stephen Eubank (TSA-2), Paula Stretz (TSA-5), Keith Bisset (TSA-5), Chris Barrett (TSA-2), and LaRon Smith (TSA-DO). For more information send e-mail to Ron Smith, Project Coordinator, at llsmith@lanl.gov, or see the Web site <http://transims.tsasa.lanl.gov/>. See the article in this issue's High-Performance Computing section.

About the front cover: Overlaid with scenes familiar to travelers, the cover shows the Portland, Oregon, case study in the center graphic, and graphics from the TRANSIMS Output Visualizer. The graphic on the right shows a screen capture of snapshot data in which the symbols represent vehicles and traffic signals. The lower center graphic depicts cumulative plan data for major streets in Portland.

Produced by the Computing, Information, and Communications (CIC) Division
 Design: Kelly Parker and Julie Medina (CIC-1)
 Managing Editor: Denise Sessions, denise@lanl.gov (CIC-1)
 Illustration: Dave Delano, Art/Illustration
 Printing: Imaging Services Group (CIC-9)
 HTML: David Van Etten (CIC-1)

BITS Contributors' Board: Hal Marshall, CIC-6; Kimberlyn Mousseau, CIC-15; Gina Fisk, CIC-5; Kathleen Jackson, CIC-3; Kei Davis, CIC-19; Nikki Gaedecke, CIC-6; David Van Etten, CIC-6; Denise Sessions, CIC-1/CIC-6; and Don Willerton, CIC-DO.

Attention Subscribers: There are two ways you can update your subscription data. E-mail bitsupdate@lanl.gov, or fill out the form inside this issue. Internal LANL subscribers please note that until we have a mechanism in place to verify our subscription database with the LANL Employee Information System we will rely on you to update your mailstop. Thanks for keeping us updated on subscription data.

Table of Contents

High-Performance Computing	2
Simplifying Performance on Clusters of Shared-Memory Multiprocessor Computers	2
TRANSIMS: Applying High-Performance Computing to Travel Forecasting	5
Information Systems	8
Electronic Journals through ABI/INFORM Database	8
New Interface for DOE Energy Science and Technology at LANL	8
Nuclear Science Abstracts at LANL Now Available	9
Desktop Computing	10
Producing Online Help for All Platforms	10
Infrastructure	14
Big Network Pipes for the ASCI Program	14
What's Happening	16
CIC-15: Sharing Information Management Expertise and Collaborating throughout the Laboratory	16
Los Alamos Hosts Interlab 2000 Conference	18
LANL Task Force on Enhancing Experimental Science	19
Research Library Resources	19
Listserv for Usability and Online Help	19
Technical and Advanced Computer Training	20
BITS Subscription Form	24
1999-2000 12-Month Index	25

Simplifying Performance on Clusters of Shared-Memory Multiprocessor Computers

by Richard Barrett, Technical Staff Member, Computational Science Methods Group, Applied Physics Division

Introduction

Large-scale physics simulations must be designed to execute effectively in a distributed memory, parallel processing environment. UPS, an acronym for "Unified Parallel Software," is a library of routines designed to help the application developer create efficient, portable, extensible, and robust large-scale parallel programs for such applications. Some parallel programming models attempt to hide the parallelism from the application writer. Others require that the application writer work at the lowest levels. UPS falls in between: it is designed to expose the parallel environment to the programmer while abstracting away the necessary complexities. The result is a simplified coding style, natural to the application, which minimizes the time spent moving data among the distributed processes.

UPS is designed to run in any computing environment that supports the C programming language (Note: UPS provides interfaces for Fortran as well as C) and which provides a method for moving data between parallel processes, such as MPI. However, the higher level interface between the application and the computing environment allows for

performance optimizations based on available hardware and software characteristics and components. In this article we give a brief overview of UPS, then describe some specific performance enhancements made to UPS for execution on the ASCI Blue Mountain compute platform. These enhancements have resulted in significant improvements in the performance of applications that are based on UPS.

Multiple SMP execution also enjoys the onbox performance benefits, but more significantly sends a single message between any two boxes rather than several. Therefore the difference between UPS and MPI remains constant regardless of message size.

Functionality

The basis of UPS is the collective view of parallel computations. The movement of data between distributed processes, I/O, linear solvers, and nonlinear solvers involve the participation of groups of processes. With advance information of certain requirements either provided by the user or ascertained during runtime, UPS can initialize, set up, and determine good ways for accomplishing the required work.

UPS functionality is distributed among components, or packages. Some of the components are briefly described below.

Basic communication enhancements.

Reductions, broadcasts, and barriers are examples of basic parallel computing functionality that has been tailored to our user's particular needs. Special capabilities have been added with regard to the operations involving large amounts of data. See Fig.1 for a performance graph of the UPS-enhanced broadcast.

Gather/scatter.

When a physics simulation operates across an unstructured mesh, the number of processes that a given process must communicate typically differs from the requirements of other processes.

Further, if the mesh changes throughout problem execution (via refinement, relaxation, or reconnection), the communication requirements also change. Keeping track of such things distracts the physicist from his or her task. UPS abstracts these details to a level which is comfortable to the programmer while at the same time to a level that allows UPS developers to optimize performance for the particular computing environment.

Data parallel operations. For convenience of programming as well as the possibility of enhanced performance, some operations on vectors distributed across the parallel processes are provided. These include inner products, vector norms, and sorts.

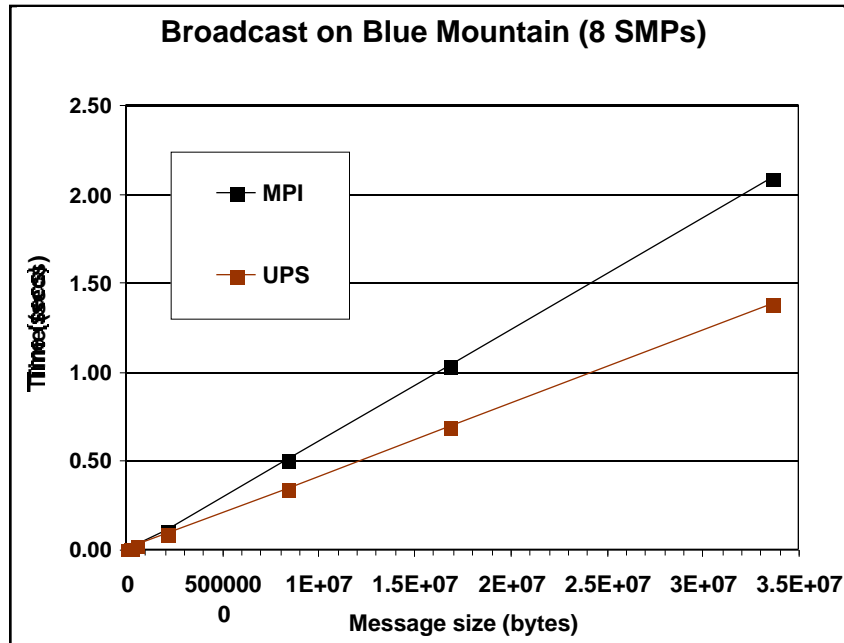


Fig 1. Performance graph of the UPS-enhanced broadcast.

Linear solvers. Many UPS users write physics simulations that require the solution of linear systems of equations. These systems are typically on a scale requiring the use of thousands of parallel processors. The solution of such systems requires sophisticated, complex mathematical and computer science technologies. This appears to conflict with the desire of the application programmer to be provided with an easy means of interfacing to such technologies. The UPS solver component, named Maya, addresses both requirements through the use of a simplified programming interface layered upon accurate and efficient solver technology. Maya is based upon Aztec, a high-quality Krylov solver package implemented and supported by Sandia National Laboratories. Local modifications include customized stopping criteria, NUMA-awareness, and support for additional matrix storage formats. Special preconditioning capabilities are being added. Interfaces to the evolving solver interface standard (ESI) are also being added.

Future functionality. UPS is an active project. Future plans include an I/O component (based upon MPI-I/O and PSF), nonlinear solvers based on the NITSOL package from the University of Utah, and adding algebraic multigrid capabilities to Maya, based upon work from the Center for Applied Scientific Computing at Lawrence Livermore National Laboratory. Ongoing work includes performance enhancements and resource efficiency issues for Blue Mountain.

Flow Control

On Blue Mountain [as well as other clusters of symmetric-memory multiprocessors (SMPs)], moving data between processes that reside on different SMPs is necessarily more expensive than moving data between processes that reside within the same SMP. Therefore, managing the flow of data between SMPs can result in a significant reduction in time spent in the communication phase. Managing the flow of data within the SMP, while

not as crucial, can also result in meaningful performance gains. In this section we describe the implementation of a particular UPS function on Blue Mountain, and show results as it is used by Partisn, an S_N neutronics application at Los Alamos. (This function is also used in the UPS implementation of sPPM.)

Partisn divides the domain into (2- or 3-dimensional) blocks. At regular intervals, data from each block must be sent to the block forward, to the right of, and below (3D only) it. This creates "waves" which sweep through the domain. This collective movement of data by the processes provides an opportunity for using the global address space (in other words, distributed shared memory) available on Blue Mountain for speeding up data movement within an SMP and between SMPs (the latter through the aggregation of messages for communication between any two SMPs). We've created the function UPS_CM_Sendrecv for this purpose.

UPS_CM_Sendrecv sets up a database for each direction in which data moves. The database includes memory locations where data will be stored and loaded (corresponding to destination and source processes) and offsets (corresponding to message lengths) into those memory locations. A "handle" is assigned to each database, which is included in subsequent calls to UPS_CM_Sendrecv. [Note: The initialization step is automatically performed whenever new parameters (destination and source processes or message lengths) are input. The user signals UPS that this is the case by inputting an undefined handle.] Now when the function is invoked by a process, the input data is moved to a predetermined location in the global address space, if necessary sent to another SMP, and when available, copied out by the target process.

Figure 2 illustrates the movement of data between distributed processes as required by Partisn.

This figure illustrates the path used for moving data between processes using UPS_CM_Sendrecv on ASCI Blue Mountain. The large boxes represent an individual SMP, while the boxes within them represent processes within the SMP. The outer circles represent the shared memory area within each SMP. The circle within the shared memory area is the buffer used to aggregate data for moving between SMPs.

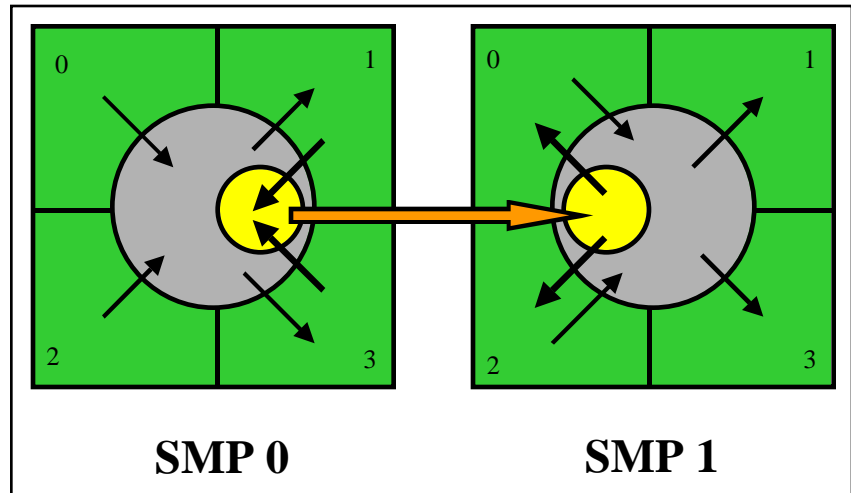


Fig. 2. Flow control illustration.

The figure on the left compares the performance of UPS_CM_Sendrecv with MPI_Sendrecv within a single SMP. The figure on the right compares the performance of the two functions when running across multiple SMPs.

Figure 3 shows the performance of this approach in comparison with the pure message passing approach previously used by Partisn.

On a single SMP, UPS outperforms MPI simply because the predetermined cooperation between processes avoids the overhead message passing that necessarily incurs. For medium size messages, UPS is about three times faster than MPI. As message size increases, MPI latency is amortized across the data transmission time, and thus MPI and UPS performance becomes comparable. Multiple SMP execution also enjoys the onbox performance benefits, but more significantly sends a single message between any two boxes rather than several. Therefore the difference between UPS and MPI remains constant regardless of message size. (For this example UPS is around six to seven times faster than MPI.) Note that speedup is constrained, however, by synchronization requirements among those processes that are sending to another SMP.

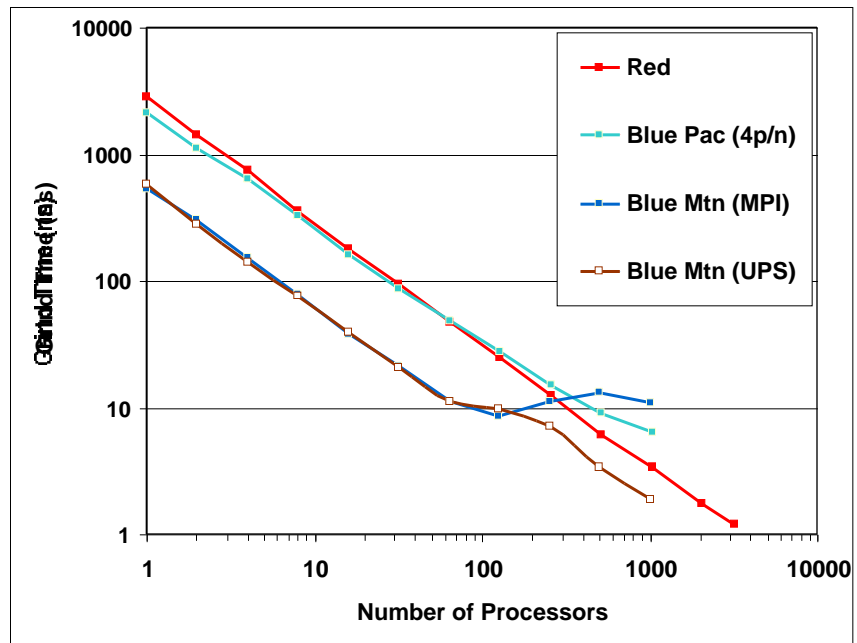


Fig. 3. Performance of UPS_CM_Sendrecv on ASCI Blue Mountain.

The current implementation of UPS_CM_Sendrecv uses MPI point-to-point (nonblocking) communication for sending data between SMPs. Future plans call for the use of a locally developed library which allows for the direct use of the HIPPI .

Conclusion

For details regarding UPS, see www.xdiv.lanl.gov/XCI/PROJECTS/UPS, or contact the UPS team at ups-team@lanl.gov.

TRANSIMS: Applying High-Performance Computing to Travel Forecasting

by Denise Sessions, BITS Managing Editor, with TRANSIMS Team Members Stephen Eubank, Basic Applied Simulation Science (TSA-2), and Kathryn Berkbigler, Computer Research & Applications (CIC-3)

Background

An interdivisional Laboratory team is engaged in a multiyear effort to develop a new transportation simulation tool that will help state and metropolitan planning organizations plagued with environmental and traffic congestion concerns. This team from TSA, CIC, and other Laboratory divisions is developing an advanced simulation and analysis system, Transportation Analysis and Simulation System (TRANSIMS), which will allow these organizations to respond to the issues raised by the U.S. Congress in the Intermodal Surface Transportation Efficiency Act and the Clean Air Act Amendments in the early 1990s. The Clean Air Act Amendments of 1990 mandate that before building any new highway infrastructure, metropolitan areas must assess the environment impact of the proposed infrastructure change.

TRANSIMS consists of mutually supporting simulations, models, and databases that use advanced computational and analytical techniques to create an environment for integrated regional transportation system analysis. By applying advanced technologies and methods, this system simulates the dynamic details that contribute to today's and tomorrow's complex transportation issues.

Tools for Decision Makers

In a recent BITS interview, TRANSIMS team member Stephen Eubank of Basic Applied Simulation Science (TSA-2) said that the team hopes to demonstrate the Laboratory's capabilities for applying simulation and modeling science to real problems. "The most pressing goal right now," according to Eubank, "is to hand off a functioning piece of software to metropolitan planning groups and let those users apply the tool to real problems."

Having worked extensively on the Traffic Microsimulator module and Selector technology, Eubank is excited about providing planners with a tool to use their data effectively because it is expensive for cities to collect data. He views TRANSIMS as a dynamical system that incorporates components of game theory, nonlinear processes, and high-performance computing.

The creators of TRANSIMS plan to put the software tools into the hands of metropolitan transportation planning organizations in fall 2001. The Laboratory is in the process of selecting a vendor to commercialize this technology and develop an end-user interface. The plan is to implement TRANSIMS for several early deployment transportation-planning organizations and establish a business infrastructure to facilitate the long-term application of the TRANSIMS technologies in the transportation-planning market.

The team has tested software versions in various case studies, such as a 25-square-mile portion of the Dallas/Fort Worth region used for demonstrating the first TRANSIMS version. Using TRANSIMS Version 1.1, the team is currently preparing to simulate the metropolitan region of Portland, Oregon, a model that requires 120,000 links and 1.5 million travelers—an order of magnitude larger than the Dallas/Fort Worth simulation of 10,000 links and 200,000 travelers. The Portland simulation also adds studies of mass transit. For more information about these case studies and other publications, see the TRANSIMS Web site <http://transims.tsasa.lanl.gov/>.

Using LANL's High-Performance Computing

Another TRANSIMS team member, Kathryn Berkbigler of Computer Research & Applications (CIC-3), who has worked on TRANSIMS planning and programming since 1993, talked to BITS about the computing aspect of the project. Berkbigler explains that all the modules in the TRANSIMS framework are object oriented, and several subsystems, particularly the network representation, are used by multiple components. (See the figure.) "It took planning at the beginning of the project to develop an extensible design so that the modules could be incrementally developed and enhanced without having to modify the entire system each time one module was improved," Berkbigler relates. All of the modules have evolved over time, adding realism and more sophisticated algorithms as the researchers gained knowledge about the simulated system. The current TRANSIMS source code contains somewhat more than 175,000 lines of code.

The first version of TRANSIMS ran on a local area network of workstations. During development, the software has been run on the Advanced Computing Laboratory's (ACL's) Nirvana supercomputer—an SGI Origin 2000 parallel computer—and more recently on the ACL's Linux clusters¹—Little Blue Penguin and Rockhopper. The current version uses 32 processors. Eubank anticipates that they could use 128–256 processors for the Portland simulation. Also, a 14-processor Sun Enterprise system is often used. See the TRANSIMS architecture in the figure.

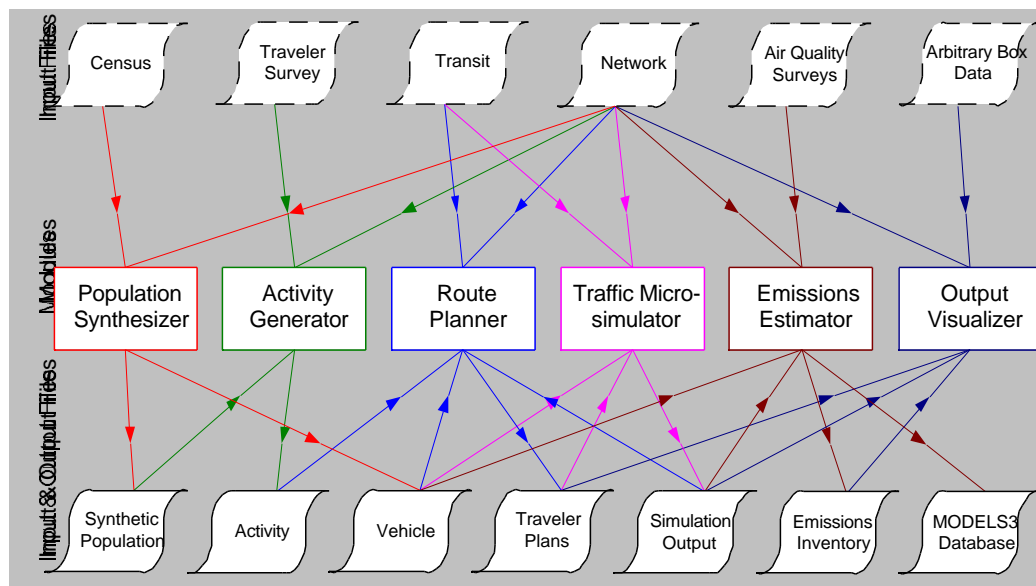
According to Berkbigler, “Running on a variety of platforms is a computing challenge that interests many developers on the project.” Simulation research as well as computational performance drives the use of high-performance computing, but metropolitan transportation planners also are expected to have access to multiprocessor machines in the near future.

How TRANSIMS Works

The system begins by building a synthetic traveler population derived from available census data. The Activity Generator module gives each traveler a set of activities based on additional survey data generated from travel diaries filled out by members of a selected sample of local households. The system uses this survey data as a template to make a daily activity list for each synthetic traveler. The Intermodal Route Planner, which develops plans for each traveler, determines the routes for the travel modes the travelers will use to traverse the transportation network to accomplish their activities. The Traffic Microsimulator then simulates the planned trips. Summary output from the traffic simulation results may be used to estimate vehicle emissions.

Accounting for human behavior is accomplished through feedback and the use of iterations and multiple loops. For example, just as people change their plans when their experiences don't meet their expectations, the TRANSIMS methods adjust the simulated travelers' behaviors to account for differences between their expectations and their experiences in the simulated real world. This adjustment is accomplished by feedback.

For each stretch of roadway and for time intervals of several minutes, the Microsimulation accumulates average travel times for use in the Activity Generator and the Intermodal Route Planner. Selectively choosing travelers whose expectations and experiences differ significantly, TRANSIMS recomputes their routes, departure times, travel modes, activity locations, or even their complete activity list. Then, the Microsimulation is executed again.



The TRANSIMS architecture from the perspective of data flow. The major TRANSIMS modules are represented in the middle row as boxes. Each of the TRANSIMS modules depends on external data that are shown in the top row. Data produced by the modules, depicted along the bottom row, are used as input to other modules.

This process continues until the travelers' expectations and experiences are approximately equivalent. Other selection criteria and halting criteria can be used in the iteration process to simulate different transportation system models, such as models of congestion management policies and their effects.

Each module is a separate executable, and all modules are typically run sequentially under the control of scripts that manage the feedback process. The Route Planner is threaded, which works well on a shared-memory machine. The Traffic Microsimulator module is an object-oriented cellular automata simulation. With cells of 7.5 m, an update step approximates a 1-second time step.

The Traffic Microsimulator uses distributed parallel execution, with a master/slave configuration. Each slave processor is responsible for simulating a portion of the total geographic region being studied, while the master synchronizes information among the slave processors. Information about the state of the system at the boundaries of the subnetwork on each processor must be communicated to adjacent processors. This information exchange must then be synchronized to correctly model physical reality. Because of differences in the number of vehicles and links assigned to each processor at any given time, the initial transportation network partition among the processors does not necessarily result in a uniform computational load across processors. Accordingly, the Traffic Microsimulator processors complete update steps at irregular intervals; ideally the developers would like to rebalance the processor computational load after 600-900 time steps.

Microsimulator Input/Output: Big Bottleneck for Parallel or Distributed Computing

During case studies, all parts of the simulation must run as fast as possible. An efficient and yet portable implementation of parallel I/O, which is critical to this goal, has been the subject of significant effort in TRANSIMS. The developers estimate that if all potential Microsimulator output is collected, each Portland run may generate about 1 terabyte of data. This is more information than developers or users want to analyze routinely, so a variety of online filtering mechanisms have been developed that may be used selectively to reduce the collected output. Output may be filtered by time so that data are gathered less frequently than every second and filtered by transportation link to collect data only in areas of interest. Data also may be collected to record when interesting events occur. Transportation planners often are interested in measures of effectiveness (MOEs) that summarize in a few important numbers the effects of alternative transportation infrastructure and policy changes. TRANSIMS provides these MOEs not only as mean values but also as distributions across the population of travelers so that the effects of transportation changes on various subpopulations can be evaluated.

Conclusion

In summary, the TRANSIMS development team has faced many challenges while creating an advanced transportation analysis tool founded in computer and simulation science and applying it to real world transportation problems that entail enormous amounts of data. They have struggled with the design of a computational simulation system involving multiple time scales, multiple modules, numerous data interfaces, and methods to select, manipulate, and move information between the modules. Concerned about potential computer hardware limitations of the ultimate TRANSIMS customers, they have researched parallel and distributed computing techniques on local area networks of workstations as well as massively parallel high-performance multiprocessor systems that handle larger and larger simulations in faster times. Finally, they have developed output collection and display interfaces that allow the users to glean the most pertinent information from the mass of data generated in a transportation system simulation.

For More Information

For more information about TRANSIMS, send e-mail to the project coordinator, LaRon Smith, at llsmith@lanl.gov.

¹For more information about the ACL's Linux clusters, see the article "Rockhopper: Extreme Linux at the ACL" and the associated miniarticle "The History of ACL's Research with Extreme Linux Clusters" in BITS Feb/March 2000 issue.

Electronic Journals through ABI/INFORM Database



Jeane Strub, Research Library, CIC-14

Eight hundred journals are now available in full text through the ABI/Inform Global database, and they have been added to the Research Library's Online Catalog (<http://lib-www.lanl.gov/cgi-bin/edb.pl?10191739>) and electronic journal Web pages (<http://lib-www.lanl.gov/cgi-bin/ejrnlsrch.cgi>). If you are looking for a specific title, when you check the Library's Online Catalog or electronic journal

pages for these titles, you will be linked directly to that journal at ABI/Inform. Examples of journals that can be found are Sloan Management Review, Cost Engineering, Human Resource Management Journal, and Personnel Management. ABI/Inform (<http://lib-www.lanl.gov/edb/umi/>) provides in-depth indexing coverage of business conditions, trends, corporate strategies and tactics, management techniques, competitive and product information, and a wide variety of other topics from 1971 to the present.

New Interface for DOE Energy Science and Technology at LANL



Irma Holtkamp, Research Library, CIC-14

A new interface and search engine is now available for DOE Energy Science and Technology at LANL at <http://doeenergy.lanl.gov/lanl/>. Features include:

- Alerts: weekly e-mail delivery of new information for searches of your choice
- Easy limiting by date, language, document type, and electronic format
- Phrase searching
- Sorting of results by relevancy ranking, date, author, or title
- Ability to search individual author names with initials
- Links to both print and electronic research library holdings
- Ability to Mark All for printing, e-mail, or downloading

DOE Energy Science and Technology at LANL contains worldwide references to basic and applied scientific and technical research literature from the U.S. Department of Energy and its contractors, other government agencies, and International Atomic Energy Agency sources. Approximately half of the references are from sources outside the United States; approximately 50% are to journal literature and 25% to technical report literature. The database has almost 4 million records covering the literature from 1974 to the present and is updated twice a month.

With this switch for DOE Energy, all of the "at LANL" indexing databases the Research Library offers now have the same basic interface and features.

Nuclear Science Abstracts at LANL Now Available

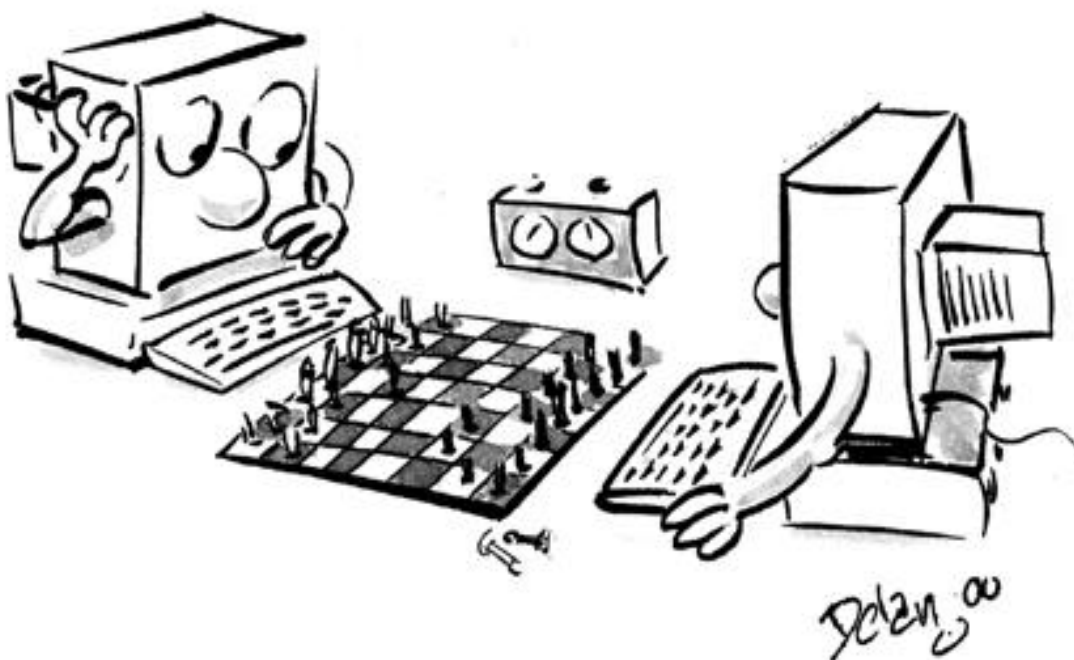
Kathy Varjabedian, Research Library, CIC-14

Most online indexes and abstracts start in the mid-1970s. Do you need research material older than that? Nuclear Science Abstracts at LANL (<http://nuclsci.lanl.gov/lanl>) is now available—a comprehensive index to historic international nuclear science and technology literature from 1948 through 1976.

The database of almost one million records includes reports of the U.S. Atomic Energy Commission, U.S. Energy Research and Development Administration and its contractors, other agencies, universities, and industrial and research organizations. This is the electronic equivalent of the printed Nuclear Science Abstracts, 1948–1976. DOE Energy Science and Technology at LANL provides coverage of the literature since 1976. About 50% of the references are to journal literature, 25% to technical report literature, and 25% to books, conferences, dissertations, and patents.

Nuclear Science Abstracts at LANL has the same search engine and interface familiar to users of the other “at LANL” databases available from the Research Library, with features such as alerts, limiting, sorting, and downloading of search results, and links to library holdings.

Send comments or questions about this new database to db-info@lanl.gov.



Producing Online Help for All Platforms

*by Marv Wetovsky & Sheila Molony,
Technical Writers/Editors,
Communication Arts and Services,
CIC-1*

With the proliferation of Web-based applications over the last several years, writers of online help have eagerly awaited integrated software tools that would take online help content and present it across all platforms. This past fall, a number of help-authoring products upgraded their packages to produce HTML-based help files.

Many products are now following the standard for online help set by Microsoft. A basic Microsoft online help system will have a table of contents, an index, and a full-text search capability. Non-Microsoft products often follow this pattern because most independent applications have first been developed for the PC, which is dominated by Microsoft products. Usability is increased by familiarity, so when users are accustomed to seeing certain navigation tools, they can find things faster. We see this now with the common use of table of contents links on the left side of a Web page.

However, as more products enter the market, the look of online help has evolved and we are now seeing variations on a theme. The table of contents might be part of a tabbed window or a second browser window that opens. Depending upon which application you use to produce your online help, the system will have a certain look.

Tools Available

Several tools are available that automate creation of the Contents, Index, and Find functions and display a help file across PC, Mac, Linux, and Unix platforms by using a browser. These would include Web-based integrated tools such as RoboHELP and ForeHelp, and HTML Web design tools such as Dreamweaver and Frontpage. Another format, JavaHelp, is cross-platform and uses its own viewer, so the user does not have to open a browser or even have one installed.

These tools work by presenting the writer a space to organize, write, and format the content, and then offering a variety of output choices. They can produce HTML files for viewing with a browser, HTML/Java files for JavaHelp, help files for use on Windows only with Internet Explorer (a nod to their history of only producing Windows files), and printed documentation (files you can open and edit using a word processor).

In addition to these integrated tools specifically developed for producing online help, you can always use a Web development tool, depending on what features you would like to have on your help system and what you want it to look like. Here you will see a brief example of these capabilities.

RoboHELP and ForeHelp

A basic help system will have an introductory page and a table of contents. Depending upon which application you use to produce your online help, the system will have a certain look. ForeHelp and RoboHELP produce help systems with a pane on the left side of the window that is divided into three tabbed sections: Contents, Index, and Search.

These help-authoring applications are two of the more popular among online help authors. You can create numerous help subjects using the application's WYSIWYG editor, add background color and images, and even edit the HTML if you wish. Both applications also generate a linked style sheet to facilitate systemwide format changes. The user interfaces, as you may expect, have a somewhat different look due to the varying design aesthetics between the two companies. In spite of this, they perform virtually identical tasks and they can each generate the same sort of output, be it HTML-based help or JavaHelp.

Figures 1 and 2 illustrate the fundamental differences between the applications' toolbars: they employ different graphical buttons and drop-down menus to accomplish tasks such as formatting text and testing the project.



Fig. 1. ForeHelp toolbar.



Fig. 2. RoboHELP toolbar.

And in Figs. 3 and 4 we see what the resulting HTML files look like when we view our projects through a browser. They are similar, but not identical, in appearance. For example, note that the ForeHelp project adds an extra icon-driven navigation bar above the contents pane.

These online help projects also feature an index and a full-text search, usually with an index pane on the left. Entering a word in the blank field at the top of the pane causes that word to become highlighted in the index listing below. Clicking the Display button displays the Topics Found window with a listing of all subjects in which the selected word is used. You can then double-click the subject of your choice to display it in the pane on the right.

Whereas the index search will search for a word that appears in the index, the full-text search option will search the entire help project to locate subjects using the search word. In this case, entering the search word displays a list of the subjects using that word. You can then double-click the subject of your choice to display it in the pane on the right.

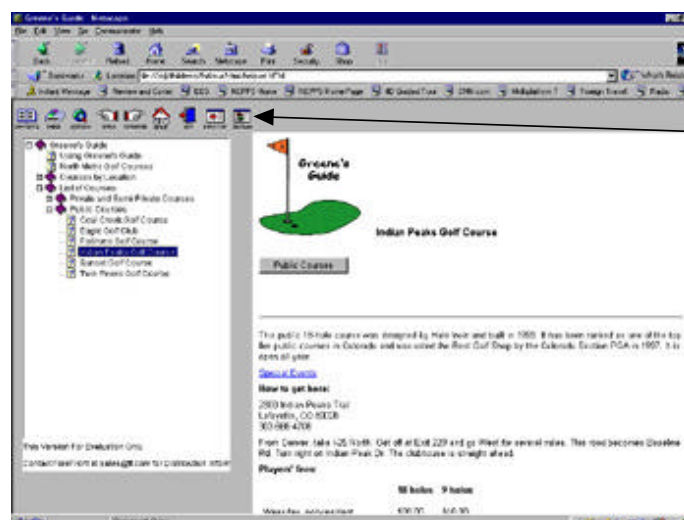


Fig. 3. ForeHelp project viewed with browser.

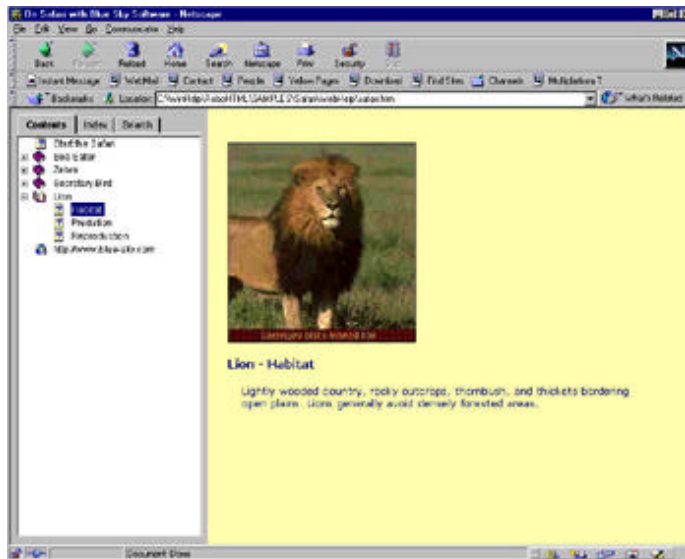


Fig. 4. RoboHELP project viewed with browser.

Web Development Tools

You may desire a different approach, however, and programs such as Dreamweaver and Frontpage have many bells and whistles absent from online help editors. This is because they are designed less with online help in mind and more toward Web site development. You have more control as to the look of your help pages and, in the case of Dreamweaver, there is a built-in FTP window for easy updates to individual pages.

JavaHelp

JavaHelp is a platform-independent format that features its own viewer, so the user does not have to open a browser. However, JavaHelp must be installed on the computer and is available as a free download from Sun Microsystems. The JavaHelp interface displays with a two-frame window. The frame on the left includes a Contents, Index, and Search tab, as shown in Fig. 5. RoboHelp and ForeHelp both have the ability to generate a JavaHelp system.

JavaHelp is also desirable because no matter which help-authoring application you use to write your help, the resulting generated JavaHelp system will look identical. Figures 5 and 6 show our two help systems, one written with RoboHELP and the other with ForeHelp. However, when JavaHelp is generated, the resulting help systems look the same, as shown in the following figures. Note that the icons at the top of the left pane represent the table of contents, index, and full text search.

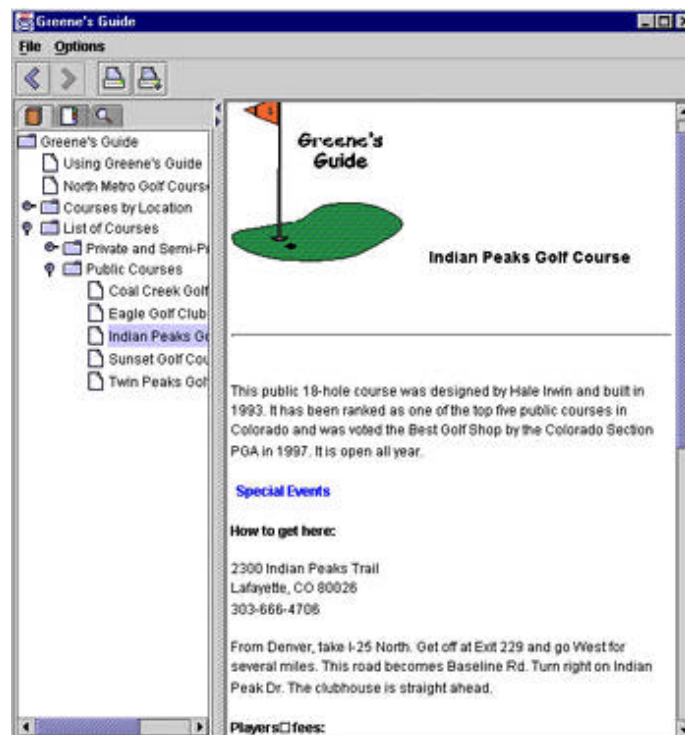


Fig. 5. User's view of ForeHelp-generated JavaHelp.



Fig. 6. User's view of RoboHELP-generated JavaHelp.

Dreamweaver is a favorite tool among the Lab's Communication Arts & Services (CIC-1) Web site developers and supplies some interesting features, such as the Behaviors Inspector, which can control the automation of different actions that occur when the user performs an action. For example, you can display a message when the user clicks on a link. This is accomplished by automatically inserting Javascript to control a given behavior. This feature saves you the trouble of having to learn and write pages of complex coding.

Dreamweaver also utilizes Javascript to create image swaps, that is, situations where an image will switch to another image when you place your mouse pointer over the original image. To see an example of a Dreamweaver help system that uses an image swap, point your browser to <http://int.lanl.gov/orgs/cic/cic1/testsite/marv/acme/testframe.html> [Editor's Note: This Web site is located inside the LANL firewall.]

and then move your mouse pointer over the menu items or Wile E. Coyote and note how the images change as you do this. Then view the source code and you'll see the Javascript that Dreamweaver has inserted, saving you a lot of time you would otherwise spend writing, debugging, and pulling out your hair. While you're at it, click some of the menu items to explore the site further.

Consistency of Appearance

Regardless of which help-authoring tool you decide to use, each can help you maintain a consistent look across your entire site. A consistent design has been shown to be an important aspect of any online site in that it projects an image of focus and cohesion, and eases the user's navigation within the site.

We can see this "standardization" being promoted here at the Lab by the use of the Laboratory Web templates and images that are downloadable from <http://www.lanl.gov/resources/templates2000/index.html>. You are encouraged to experiment with these template images and use them as you see fit. Any of the tools discussed here can be used to insert a template image and place it where you want it on the page. Pick the one that best suits your needs.

Examples

We currently have several online help projects in various stages of development, including one for the LANL Foreign Travel System for which we are using Dreamweaver. To view it, point your browser to <http://truchas.lanl.gov/foreigntravel/entry.html>. Then click the word "Help" at the bottom of the list on the left to open the help window.

If you'd like to see some other online help examples we have under development using some of the other tools discussed here, please feel free to contact Marv (marv@lanl.gov) or Sheila (sheila.molony@lanl.gov) and we'll be happy to show you around and discuss possibilities for any applications you may have in development.

Big Network Pipes for the ASCI Program

by Stephen Tenbrink, Deputy Group Leader, CIC-5 Network Engineering

The Accelerated Strategic Computing Initiative (ASCI) program, which has deployed new high-end computing systems at the three DOE Defense Programs (DP) laboratories, has another component requiring higher bandwidth between these laboratories. With the decision to place a 10-teraflop "ASCI White" computer at Livermore at the end of this year and a 30-teraflop computing system at Los Alamos at the end of 2001, there will be an imbalance of compute power among the three DP laboratories. The decision to create this imbalance was based on the need to avoid the high costs required to deploy these systems at all three sites. Along with this decision it was also apparent that the other two laboratories that did not have the current high-end ASCI system would need access to the machine and that this access should "appear as if local to the extent possible" to the remote user.

ASCI's Distance Computing

Out of this was born the "Distance and Distributed Computing" program also known as DisCom. The first phase of DisCom is to address the distance computing aspects of ASCI given the imbalance of computing power mentioned above. The second phase is to build upon the first phase by creating a wide area network (WAN)

computing fabric for the entire DOE Nuclear Weapons Complex that includes distributed computing resources located throughout the complex. Sandia National Laboratory is the lead lab for DisCom but both Los Alamos and Livermore have significant involvement in this effort.

So, what does it mean for a remote user to gain access that appears to the user "as if local..."? To respond to this requirement each lab interviewed some of their users to see what operations were needed to get meaningful results in a timely manner. (Note the words "timely manner" in the last sentence.) These operations involve data movement, as would be expected, but more important it's what is done with the data at the remote site that is critical. This could involve local archival storage, intermediate storage, high-resolution visualization with various forms of rendering, etc. After the interviews and comparing results from the other labs, the DisCom team found the answer was "all of the above." In other words, DisCom will need to create a WAN that interconnects the network backbones at each of the three laboratories (and an adjunct link to Sandia/California) with enough bandwidth that will support any of the network functions that users now do locally. To achieve this, very high-speed network data pipes would be needed. The only problem that DisCom will not be able to overcome is the latency issue due to the "time of flight" of the data from one site to another.

This could be tens of milliseconds if the user is in California and the ASCI computer is at Los Alamos. This is why the phrase "to the extent possible" was added to the DisCom goal.

How Big the Bandwidth?

DisCom project personnel then determined the size of the data pipes. This was done several ways but the guiding principal was to move "x" terabytes of data between sites in about an hour where "x" was the expected size of the problem data sets anticipated from the current and future ASCI high-end machines. While this is not an exact mathematical approach it would give us a good estimate of what bandwidth would be needed. The result was the approach showed that the minimum bandwidth should near the teraflop rating of the machine in gigabits/sec. Thus, for the 10-teraflop Livermore White computer, a 10-gigabit/sec link should suffice; for the Los Alamos 30-teraflop, a 30-gigabit/sec link; and for the future 100-teraflop ASCI machine, a 100-gigabit/sec link. When compared to the current interlab 155 megabits/sec WANlink (provided by ESNet) that LANL has, these rates seem astounding. But given the size of the data sets expected from the future ASCI machines and the goal of creating an "as if local" environment, these rates would be required.

The Telcos' Response to Future Growth

Luckily, the growth of the Internet is helping DisCom reach these bandwidth goals. Because of the high demand for digital communications for Internet activity, most major long distance telecommunication company carriers (Telcos) are installing new fiber optic links around the country and, as they install these links, they are adding extra capacity (addition fiber optic strands) for future growth. Accompanying this is a new fiber optic technology called Wave Division Multiplexing (WDM) and Dense Wave Division Multiplexing (DWDM) which allow one fiber strand to carry multiple data streams on different wavelengths (colors). These two efforts are combining to drive down the cost of wide area bandwidth to a point that even 100 gigabits/sec between California and New Mexico is within the budgetary constraints of DisCom. Two years ago it was thought that the Telcos would have trouble even providing 100 gigabit/sec. Today most of the Telcos we've approached hardly blink an eye when we state our requirements.

LANL, however has a unique problem in this area. Most long distance carriers have a "point of presence" (POP) in major metropolitan areas and rely on the local exchange carrier (LEC) to go the "last mile" from the POP to the final destination. Being that Los Alamos is where it is, this becomes a "last 100 mile" problem for the Laboratory. Our LEC, US West, provides the communication infrastructure for this last 100 miles that supports your long distance phone service and the Laboratory's ESNet link to a specific POP in Albuquerque depending on who the long distance carrier is. The problem is that we do not know if US West has the fiber capacity to support 30 or 100 gigabit/sec, even with DWDM. The answer to this question will be determined in the next few weeks when

responses to an Request For Quotation (RFQ) for the DisCom WAN are evaluated by DisCom TriLab members. Luckily, again because of the growing demand for Internet access, other companies are interested in building fiber links into northern New Mexico. These companies are also bidding on the DisCom RFQ. It is hoped that the result of all of this will be a WAN that provides the required DisCom bandwidth and, in addition, provides the infrastructure to support increased bandwidth into northern New Mexico to meet the growing demand for Internet access, both commercial and residential.

For more information about the TriLab DisCom project, see the DisCom Web site at <http://www.cs.sandia.gov/discom/>, or contact Steve Tenbrink at sct@lanl.gov.



CIC-15: Sharing Information Management Expertise and Collaborating throughout the Laboratory

by Denise Sessions, BITS Managing Editor, with Bethany Wannigman, Group Leader, Advanced Database and Information Technology Group, CIC-15

This article is one in a series of interviews BITS is conducting with CIC managers to get their views of the "big picture" as it relates to their work and the Laboratory mission. These people have also been asked to do a little forecasting as it applies to their business. BITS invites readers to join in the spirit of these interviews, treating the forecasts as a sort of informed speculation without holding anyone's "feet to the fire" to make the predictions come true.

Advanced Database and Information Technology Group (CIC-15) Group Leader Bethany Wannigman envisions that a critical step for information technology management at the Lab will be to define the Chief Information Officer (CIO) position, establish its location in the Lab's management structure, and appoint the CIO. In general, a CIO's primary goal is to ensure consistency in information and data management across an organization (see sidebar).

Wannigman, who has been the CIC-15 group leader beginning this January, says she thinks that at this time the Lab needs a focused approach to information management (IM). "Now more than ever," says Wannigan, "information systems are becoming more critical to the Laboratory." Her group's role is to increase cost effectiveness of IM and make information systems more efficient. For example, they have purchased an Oracle combined network license that will save the Laboratory from buying individual licenses and paying system administration fees for individual accounts. Wannigman also sees the need for new systems to interface with current Laboratory systems in order to achieve better data integration.

Wannigman sees her group integrating and collaborating with other Lab organizations involved in IM and working toward positioning the Laboratory toward a common IM "face." She participates on the CIC IM Focus Team, a Laboratory group organized to improve information services into a coherent, focused approach to IM. Wannigman and various CIC-15 group members work with the Lab's Information Architecture (IA) project and its focus teams. For example, CIC-15 members are participating in IA's Web Computing Team and the Enterprise Application Planning Team.

Fresh Impact with New CIC-15 Structure

Recently Wannigman implemented changes to the group's structure to address evolving customer needs and position the group's resources into interactive, overlapping, and multi-disciplinary teams. The group serves customers at their sites as well as develops and coordinates customer requests in an integrated way to provide the best products and services to meet customer needs. For example, the Software Engineering Process team specializes in project management. In response to a customer request, the team gathers and analyzes customer requirements and develops a proposal to meet the customer needs. To increase value to customers, CIC-15 staff minimize costs by leveraging previous software development and maximizing code reuse.

The group is organized into four teams. The Software Engineering team offers these services: project management; system requirements, analysis, and design; and quality assurance. The Transactional Development team handles structured data, whereas the Full-Text Development team works with unstructured data. The Infrastructure team provides services associated with databases such as Oracle, Access, and Sybase; system administration; and document repositories. To get an

idea of the range of CIC-15 products and services and the expertise resident in CIC-15, see their Web site: <http://iosun.lanl.gov:2001/htmls/cic15/external/p+s.html>.

Group Leader Background

Wannigman earned her MS in Business Computing Science from Texas A&M University. She started working at the Laboratory as a staff member (programmer/analyst) in Business Information Systems (CIC-13) in 1986. During her 13-year tenure in CIC-13, Wannigman worked for 11 years in roles such as project leader and team leader analyzing, developing, implementing, and testing Lab-wide systems and software, for example, the Employee Information System and Financial Management System. During the last two years in CIC-13, Wannigman served as Acting Deputy Group Leader. Her hobbies include watching her daughter play soccer, raising animals, and playing tennis and bridge.

Chief Information Officer Sidebar

Currently an effort is underway to determine what a CIO position at the Laboratory would include. The CIO Council is a separate entity from the team of people designated to address the CIO position. The CIO Team (which is a separate entity from the CIO Council¹), chaired by CIC Division Leader Charlie Slocomb, is looking at issues such as where a CIO would reside organizationally, how a CIO would be held accountable, and how the position would be funded. In considering a CIO position, the Lab is following a DOE-wide trend, as other institutions across the DOE complex have already appointed, or are looking at appointing, a CIO.

Since the CIO Team began meeting in February, it has surveyed companies to study the CIO model used by those corporations. In general, a CIO's primary goal is to ensure consistency in information and data management across an organization. In conducting its survey, however, the team found great variation in the ways CIOs function. Some CIOs report to the vice-president of research and development, other CIOs report to the vice-president in charge of finance and human resources, while yet other CIOs reside at the top level of a company's management. Charlie Slocomb presented the team's findings and recommendations to the Senior Executive Team in April 2000.

¹ For more information about the CIO Council, see the BITS article "The CIO Council—An Institutional Resource" in the April/May 2000 issue.

Los Alamos Hosts Interlab 2000 Conference

by Denise Sessions, BITS Managing Editor, with Chris Lindberg, Web Designer Programmer, Communications Arts & Services, CIC-1

October 31 through November 2, Los Alamos will host Interlab 2000—a conference designed for DOE Webmasters, Web developers, and managers of Internet. Hosted by LANL for the first time, this is the 7th annual Interlab conference.

The Interlab 2000 planning team polled potential participants about their expectations for this year's conference. The survey addressed hot topics, cost expectations, likes and dislikes (from previous Interlab conferences), and additional topics of interest. To see the complete survey results, see the Interlab 2000 planning Web site <http://bus.lanl.gov/interlab2k>. According to the 154 survey respondents, the top four topics of interest are:

- Dynamic Web Building Using Databases
- Document Management/Search Tools
- Web Graphics/Design
- Security Concerns

Tentative Agenda

The 3-day event will feature live presentations, breakout sessions and laboratory demonstrations, a vendor exhibit, poster session, two networking receptions, and four keynote presentations. Using hot topics identified in the survey results, the planning team decided on these keynote speakers.

Design keynote speaker: Myke Ninness, Adobe Systems.

Myke is a senior product manager at Adobe Systems. He has 10 years experience as an instructor, hands-on trainer, and conference presenter. He was a featured artist in Photoshop Studio Secrets and is the author of Photoshop 5 Web Magic and Photoshop Power Shortcuts.

Security keynote speaker: Matt Curtin, Interhack Corporation.

Matt is founder and original hacker at Interhack Corporation (<http://www.interhack.net/>), directing the research and development of tools, technology, and services for privacy and data security. In addition, he is a lecturer at Ohio State University's computer and information science department, teaching Common Lisp.

Database keynote speaker: David Phipps, Quios.

As Director of Engineering, David is responsible for managing the office toys (bubbles, balls, and squirt guns), as well as the Quios team of programmers. David has over 16 years experience in technology development and management. Before joining Quios, David was Director of Engineering at Zing, a photo-hosting community, where he was in charge of Web scaling, applications design, and development.

Knowledge Management keynote speaker: Scott Cooper, IBM/Lotus.

Scott is vice president of Lotus' Knowledge Management Products Group (KMPG) where he leads Lotus' entry into the highly strategic KM market segment. In his current role, Scott has been instrumental in designing the framework and direction for Lotus' knowledge management products and solutions. Before joining Lotus, Scott was a product engineer at Eastman Kodak where he headed the development and launch of Eastman Kodak's first PC-based desktop imaging product. Scott speaks extensively at industry trade shows around the world, and is frequently quoted in industry press articles.

Help Create the Conference

The Interlab 2000 planning team is offering numerous ways for participants to help create the conference and the supporting Web material. You can submit proposals for presentations at the conference, papers (for posting on the Web), and posters for the poster session. The planning team is having the first ever Interlab logo contest. The winner will receive free registration to the conference. For more information, see the logo contest rules on the conference Web site <http://i2k.lanl.gov>. The deadline for submissions is June 30, 2000.

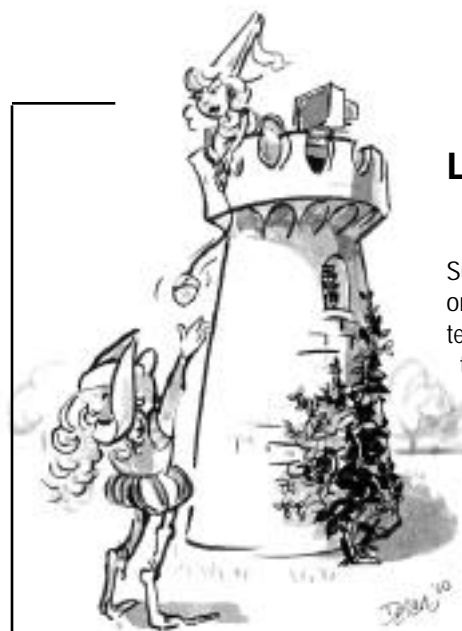
For timely information about details of the conference, and to see participants' contributions as they are added to the site, check the conference Web site often.

LANL Task Force on Enhancing Experimental Science

Lab Director John Browne and Deputy Director Bill Press have chartered the LANL Task Force on Enhancing Experimental Science to look into issues surrounding the effective conduct of experimental science at the Laboratory and provide recommended actions to the Senior Executive Team. Led by Senior Fellow Gerry Garvey, the task force is seeking information from managers and employees, from both technical and support organizations, who have ideas and views to contribute to this evaluation. Further information and opportunities for participation are presented on the task force's Web page at (<http://int.lanl.gov/taskforce/experimentalscience/>). The task force is scheduled to develop a preliminary report by June 1 and a final report in September, and will continue to accept and consider input throughout the process.



The LANL Research Library offers a variety of training opportunities for the Laboratory community. Available sessions focus on specialized library databases and other electronic resources. A complete list of course offerings can be found at <http://lib-www.lanl.gov/libinfo/training.htm>. All sessions are available to individuals or groups at the library or your site. Arrange for a session by contacting the Library, phone 7-4175 or e-mail library@lanl.gov. Library tours are available on a drop-in basis every Wednesday at 1:00 p.m.



Listserv for Usability and Online Help

Subscribe to a mailing list for people interested in software documentation and online help. It's a convenient way to refer each other to interesting articles and techniques, and perhaps help each other with specific problems from time to time.

To be added to the list, send an e-mail message to listmanager@lanl.gov with a single line in the body of the message:

subscribe documenters

After subscribing, you will receive a confirmation message that also explains how to unsubscribe.

Computer Training

The Customer Service Group (CIC-6) offers technical computer training (enterprise information applications, communications, office administration, and Web authoring) and advanced computer training (programming languages, system administration, and advanced applications). To register for a course access our Web page at <http://www.lanl.gov/internal/training/training.html>. Or from the LANL home page select the links: Training, Computer. For further information about technical computer training call (505) 667-9559, and for advanced technical computer training call (505) 667-9399.

Technical and Advanced Technical Computer Training Courses

Communications

- Eudora 4.2
- Meeting Maker 5.5.3

Office Skills 2000

- Office Skills 2000–LANL Computing

Web Authoring and Browsing

- Dreamweaver 3.0
- FrontPage 2000
- HTML

Enterprise Information Applications (EIA)

- Date Warehouse–Basics
- Date Warehouse–EDS Reports
- Data Warehouse–Passport
- EDS–Basics
- EDS–GUI
- EDS–Training Plans
- Infomaker
- Invoice Approval System
- Procurement Desktop
- Purchase Card System
- Recharge
- Time & Effort GUI
- Travel Foreign
- Travel Domestic GUI
- Web JIT

Other EIA Courses

- Automated Chemical Information System (ACIS)
- Directory Information System (DIS)
- Financial Management Information System (FMIS)
- Key/Core
- Property Accounting, Inventory, and Reporting System (PAIRS)
- Resource Planning Module (RPM)
- Salary Management Tool (SMT)
- Secretarial/Contract Service (SE)
- Signature Authority System (SAS)

System Administration Training

- Advanced Citrix
- IRIX (SGI) Network Administration
- IRIX (SGI) System Administration (Advanced)
- IRIX (SGI) System Administration (Beginning)
- Linux System & Network Administration
- Metaframe 1.8 Administration (Citrix)
- SGI Performance Evaluation and System Tuning
- Solaris 7 Network Administration II
- Solaris 7 System Administration I
- Solaris TCP/IP Network Integration
- Unix and Windows NT Integration
- Windows 2000 Classes (to be determined)
- Windows NT Optimization and Troubleshooting
- Windows NT Security
- Windows NT Workstation and Server

Technical and Advanced Technical Computer Training Courses—Continued

Programming Training

- Fortran 90/95
- Advanced Perl Programming with CGI
- ANSI/ISO C++ Programming Clinic (Advanced C++)
- C Programming (Advanced)
- C Programming (Beginning)
- C++ for Experienced C Programmers
- Distributed Objects Using Corba
- Distributed Programming with Java
- Java Programming
- Java Programming Workshop
- JavaScript: Building Interactive Web Sites
- Object-Oriented Analysis and Design
- Perl Programming
- Programming for Beginners Using Java
- Shell Programming

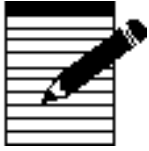
Application Training

- Apache Web Server
- C++ and the Unified Modeling Language
- Fortran – lint & lint-PLUS
- Foundations of IDL Programming
- FrameMaker Basic and Advanced
- IDL 5.0 Graphic Object Workshop
- Netscape Servers for Intranet Development
- Origin 2000 Applications Programming and Optimization
- Sendmail—Managing Internet Mail
- Sybase Performance and Tuning for System 11
- Unix (Beginning)
- Unix (Advanced)
- Visual Basic 5.0 Fundamentals
- Visual C++ Windows Programming
- Visual Fortran

Note: You do not need an ICN password to use e-mail.

What's Happening

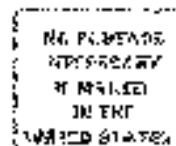
Note page . . .



Reader Feedback

Feedback helps us to provide a document that responds to the changing needs of its readership. If you have comments or questions about this publication, please let us hear from you. We have reserved the back of this form for that purpose. We also accept articles for publication that are of interest to our readers. Contact the managing editor for more information. This form is also used for new subscriptions, deletions, or changes. Instructions are on the back. If you prefer to contact us by e-mail, send your comments and/or subscription request to denise@lanl.gov.

Do Not Staple
Fold on This Line First



BUSINESS REPLY MAIL

FIRST-CLASS MAIL PERMIT NO. 1000 LOS ALAMOS, NM

POSTAGE WILL BE PAID BY THE ADDRESSEE

MAIL STOP B251
ATTN: DENISE SESSIONS, MANAGING EDITOR
CUSTOMER SERVICE GROUP (CIC-6)
LOS ALAMOS NATIONAL LABORATORY
PO BOX 1663
LOS ALAMOS, NM 87545-9916



Do Not Staple, Seal with Tape
Fold Here

cut along dashed line

Feedback

This image shows a single sheet of white paper with horizontal ruling lines. The lines are evenly spaced and run across the width of the page. There are no margins, text, or other markings on the paper.

New Subscription, Deletions, and Changes

Bits is published by Los Alamos National Laboratory. If you would like to be added to or deleted from our mailing list, please check the appropriate line, complete the form below, and mail us the form, or e-mail bitsupdate@lanl.gov

_____ Add my name to the BITS mail list.

_____ Delete my name from the BITS mailing list.

_____ Change my name/address as indicated below.

Name	Date	
Address	Mail Stop	
Group	Organization	
City	State	Zip
Phone	Number of copies	

1999–2000 12-Month Index

To electronically access any of the articles listed on this index go to http://www.lanl.gov/orgs/cic/cic6/bits/index/index_home.html. Note that the page numbers printed on this page correspond to the page numbers in the printed versions of the publication.

Keywords	Title of BITS Article	Date	Page
ASCI	<i>Tecolote: An Object-Oriented Framework for Hydrodynamics Physics</i>	Aug. 99	26
	<i>Using Tecolote Components to Extend Object-Oriented Programming</i>	Sept. 99	19
	<i>Keeping Informed about Progress on the Strategic Computing Complex Construction Project</i>	Jan. 00	19
	<i>Blanca/Tecolote Base Class Redesign</i>	Feb./Mar. 00	2
Awards	<i>BITS Publications Earns Society for Technical Communication Awards</i>	Feb./Mar. 00	28
BITS Interviews	<i>Ann Hayes</i>	J/J 99	22
	<i>CIC Division Hosts Student Poster Presentations</i>	Sept. 99	23
Book Review	<i>Philip and Alex's Guide to Web Publishing</i>	Nov. 99	26
CIC (Computing, Information, & Communications)	<i>CIC Division Undergoes Review in March</i>	Feb./Mar. 00	29
	<i>External Committee Reviews CIC Division Progress and Future Directions</i>	Apr./May 00	23
Desktop Computing	<i>Adjusting for Reality: Mitigating Uncertainty in Projects</i>	Aug. 99	7
	<i>Your Data Is Gone . . . Now What?</i>	Aug. 99	11
	<i>Desktop on Demand: A Desktop That Displays on Your Computer, but Actually Runs on the Enterprise Server</i>	Sept. 99	6
	<i>Grassroots Software Management: Simple Things to Do</i>	Sept. 99	8
	<i>News About Microsoft Networking at Los Alamos: The LANL Master Accounts Domain Is Up and Running!</i>	Sept. 99	10
	<i>Notes on Setting Up and Configuring an NT Domain</i>	Sept. 99	12
	<i>A Primer for Developing Online Help</i>	Oct. 99	4
	<i>News about Laboratory Electronic Mail and Policies</i>	Oct. 99	6
	<i>ADSM: A LANL Centralized, Automated Backup and Restore Option</i>	Nov. 99	7
	<i>Scalable Vector Graphics—Web Graphics with Original-Quality Artwork</i>	Nov. 99	2
	<i>Ghost Helps CIC-2 Back Up Data</i>	Jan. 00	11
	<i>ESD Customer Survey Results</i>	Feb./Mar. 00	17
	<i>Software Licensing Misconceptions</i>	Feb./Mar. 00	20
	<i>SMS or Windows 2000 for Desktop Management?</i>	Apr./May 00	8
	<i>Windows 2000: How Well Can the Lab Use This Tool?</i>	Apr./May 00	11
Linux	<i>Rockhopper: Extreme Linux at the ACL</i>	Feb./Mar. 00	6
	<i>The History of ACL's Research with Extreme Linux Clusters</i>	Feb./Mar. 00	11
HTML	<i>Web Content Architecture: Taming the Tangle of Protocols</i>	Sept. 99	15
	<i>MathML: A Kind of "TeX for the Web"</i>	Oct. 99	2
	<i>Scalable Vector Graphics—Web Graphics with Original-Quality Artwork</i>	Nov. 99	2

Keywords	Title of BITS Article	Date	Page
High-Performance Computing	<i>Complex Systems Modeling: Using Metaphors from Nature in Simulation and Scientific Models</i>	Nov. 99	18
	<i>Dimensionless Coding Techniques in Tecolote: Using a Single Source-Code Base for Multidimension Programs</i>	Nov. 99	22
	<i>Lecture Review: Next-Generation Chips, Processors, Transistors, and Wiring</i>	Aug. 99	23
	<i>Tecolote: An Object-Oriented Framework for Hydrodynamics Physics</i>	Aug. 99	26
	<i>Frameworks Are Models, Too!</i>	Oct. 99	16
	<i>Automated OS Install HOWTO: Linux Cluster via Gigabit Ethernet</i>	Oct. 99	20
	<i>Achieving Revolution through Evolution: Breaking Giant Leaps into Small Steps</i>	Jan. 00	2
	<i>Ocean and Climate Modeling</i>	Jan. 00	6
	<i>Rockhopper: Extreme Linux at the ACL</i>	Feb./Mar. 00	6
	<i>The History of ACL's Research with Extreme Linux Clusters</i>	Feb./Mar. 00	11
	<i>Water Resources—Applying LANL's High-Performance Computing Capability to a National Problem</i>	Apr./May 00	2
Infrastructure	<i>Electronic Information Protection Regimes</i>	Aug. 99	13
	<i>Web Content Accessibility: New W3C Guidelines Have Arrived</i>	Aug. 99	20
	<i>Web Content Architecture: Taming the Tangle of Protocols</i>	Sept. 99	15
	<i>Data Analysis in the Advanced Surveillance Technology Initiative: AMISS Testbed</i>	Oct. 99	9
	<i>Transaction Integrity and Validity in MC&A Databases</i>	Oct. 99	12
	<i>Take a Look at the New PAGES</i>	Nov. 99	11
	<i>Recharge System Telecommunications Enhancements</i>	Nov. 99	13
	<i>Fast Reduction of Analytical Data Using WINDOWS® 95/98/NT: A Software Application for ICP Mass Spectrometry</i>	Jan. 00	13
	<i>New Tool for System Administrators of LANL NT MAD Accounts: Domain Manager Application</i>	Jan. 00	16
	<i>Moving an NT Workstation from a Workgroup to a Domain</i>	Feb./Mar. 00	22
	<i>Information Architecture: Revitalizing This Old House</i>	Apr./May 00	14
Information Systems	<i>The CIO Council—An Institutional Resource</i>	Apr./May 00	16
	<i>aha!—An Interface to the Best of LANL Science & Technology Web Resources</i>	Aug. 99	2
	<i>More Than a New Look & Feel: The New LANL Badging System</i>	Aug. 99	5
	<i>Laboratory Overview Presentation Materials on the Web</i>	Sept. 99	2
	<i>New INSPEC® at LANL</i>	Sept. 99	28
	<i>Keep on Top of Information Technology with Gartner Group</i>	Sept. 99	28
	<i>Beilstein Organic Chemistry Database Is Now on the Web—Free</i>	Nov. 99	27
	<i>Books in Print on the Web</i>	Apr./May 00	6
	<i>Report Library's Classified Database</i>	Apr./May 00	6
	<i>SciSearch® at LANL/Social SciSearch® at LANL Releases New Version</i>	Apr./May 00	7

Keywords	Title of BITS Article	Date	Page
<i>Information Systems –continued</i>	<i>New Interface for Biosciences Database</i>	<i>Nov. 99</i>	<i>27</i>
	<i>New Version of Science Server[®] at LANL</i>	<i>Jan. 00</i>	<i>10</i>
	<i>New Engineering Index[®] at LANL</i>	<i>Jan. 00</i>	<i>10</i>
<i>Networks</i>	<i>News about Networking at Los Alamos: The LANL Master Accounts Domain Is Up and Running!</i>	<i>Sept. 99</i>	<i>10</i>
	<i>Notes on Setting Up and Configuring an NT Domain</i>	<i>Sept. 99</i>	<i>12</i>
	<i>New Tool for System Administrators of LANL NT MAD Accounts: Domain Manager Application</i>	<i>Jan. 00</i>	<i>16</i>
<i>Security</i>	<i>Policy Changes for Foreign Visits and Assignments</i>	<i>Jan. 00</i>	<i>21</i>
	<i>Publication Release and Accountability at the Laboratory</i>	<i>Jan. 00</i>	<i>23</i>
	<i>An Update on Computer Security Initiatives</i>	<i>Feb./Mar. 00</i>	<i>25</i>
<i>Training</i>	<i>New Options for Getting Work-Related Training</i>	<i>Apr./May 00</i>	<i>24</i>
	<i>Technical and Advanced Computer Training</i>	<i>Apr./May 00</i>	<i>26</i>
<i>Usability and Online Help</i>	<i>Listserv for Usability and Online Help</i>	<i>Apr./May 00</i>	<i>25</i>
<i>Windows NT</i>	<i>News about Networking at Los Alamos: The LANL Master Accounts Domain Is Up and Running!</i>	<i>Sept. 99</i>	<i>10</i>
	<i>Notes on Setting Up and Configuring an NT Domain</i>	<i>Sept. 99</i>	<i>12</i>
	<i>New Tool for System Administrators of LANL NT MAD Accounts: Domain Manager Application</i>	<i>Jan. 00</i>	<i>16</i>
<i>World Wide Web (WWW or Web)</i>	<i>Scalable Vector Graphics—Web Graphics with Original-Quality Artwork</i>	<i>Nov. 99</i>	<i>2</i>
	<i>Web Design 2000: A Conference for LANL Web Communicators</i>	<i>Apr./May 00</i>	<i>18</i>
<i>Writing Resources</i>	<i>Telling Your Story in BITS: Sharing Your Expertise and Energy</i>	<i>Aug. 99</i>	<i>31</i>
<i>Year 2000 (Y2K)</i>	<i>Year 2000: Are You Still Ready?</i>	<i>Nov. 99</i>	<i>14</i>

What's Happening

Note page . . .



Customer Support Center (505) 665-4444, ext. 851, or cichelp@lanl.gov

Because of a wide variety of CIC computing services, numerous facilities are available to address your questions. If you are uncertain whom to call, you can always call the Customer Support Center (CSC). CSC consultants are trained to either answer your question or locate someone who can. To reach the appropriate consultant, dial 665-4444 and make your selection from the following choices:

Option 1: New user topics including e-mail, Eudora, secure passwords, new user token card registration, and Laboratory terminations

Option 2: Enterprise Information Applications such as Travel, Time and Effort, Purchase cards, and Authorities

Option 3: Scientific computing, storage systems, and networking

Option 4: Technical computer instruction and training

Option 5: Desktop Consulting for PC and Macintosh software support.

Consulting Via E-Mail

Customer Support cichelp@lanl.gov

Scientific and engineering computing consult@lanl.gov

Administrative and business computing eiaconsult@lanl.gov

Token cards and secure passwords validate@lanl.gov

Macintosh computing Mac-help@lanl.gov

PC computing PC-help@lanl.gov

UNIX computing UNIX-help@lanl.gov

Other Useful Numbers

Advanced Computing Laboratory 665-4530

Central Computing Facility 667-4584

Network Operations Center noc@lanl.gov or 667-7423

Telephone Services Center 667-3400

Los Alamos National Laboratory, an affirmative action/equal opportunity employer, is operated by the University of California for the United States Department of Energy under contract W-7405-ENG-36.

All company names, logos, and products mentioned herein are trademarks of their respective companies. Reference to any specific company or product is not to be construed as an endorsement of said company or product by the Regents of the University of California, the United States Government, the U.S. Department of Energy, nor any of their employees. Los Alamos National Laboratory strongly supports academic freedom and a researcher's right to publish; as an institution, however, the Laboratory does not endorse the viewpoint of a publication or guarantee its technical correctness.

Los Alamos
NATIONAL LABORATORY



These artist conceptions show the future home for some employees from Telecommunications (CIC-4) and Network Engineering (CIC-5). Gerald Martin Construction of Albuquerque, the general contractor for the project, has already started construction at the site east of Diamond Drive in Technical Area 3 adjacent to the US West operations dispatch facility. The 24,000 square-foot two-story structure is scheduled to be completed in September, weather permitting.

Los Alamos

NATIONAL LABORATORY

Mailstop B251 Los Alamos, New Mexico 87545
A U.S. Department of Energy Laboratory

BITS is published bimonthly to highlight recent computing and communications activities within the Laboratory. We welcome your suggestions and contributions.

BITS may be accessed electronically at this URL:
<http://www.lanl.gov/cic/publications.html>

LALP-00-2 (6-7/00)

Nonprofit
 organization
 US Postage

PAID

Albuquerque, NM
 Permit No. 532